arXiv:1403.8018v1 [q-fin.RM] 31 Mar 2014

Are credit ratings time-homogeneous and Markov?

Pedro Lencastre[1,2], Frank Raischel[3], Pedro G. Lind[4], Tim Rogers[5]

[1] ISCTE-IUL, Av. Forças Armadas, 1649-026 Lisboa, Portugal
(e-mail: pedro.lencastre.silva@gmail.com)
[2] Mathematical Department, FCUL, University of Lisbon, 1749-016 Lisbon, Portugal
[3] Instituto Dom Luiz, University of Lisbon, 1749-016 Lisbon, Portugal
[4] ForWind and Institute of Physics, University of Oldenburg, Ammerländer Heerstrasse 136, DE-26111 Oldenburg, Germany
[5] Centre for Networks and Collective Behaviour, Department of Mathematical Sciences, University of Bath, Claverton Down, BA2 7AY, Bath, UK

Abstract. We introduce a simple approach for testing the reliability of homogeneous generators and the Markov property of the stochastic processes underlying empirical time series of credit ratings. We analyze open access data provided by Moody's and show that the validity of these assumptions - existence of a homogeneous generator and Markovianity - is not always guaranteed. Our analysis is based on a comparison between empirical transition matrices aggregated over fixed time windows and candidate transition matrices generated from measurements taken over shorter periods. Ratings are widely used in credit risk, and are a key element in risk assessment; our results provide a tool for quantifying confidence in predictions extrapolated from these time series.
Keywords: Generator matrices, Continuous Markov processes, Rating matrices, Credit Risk.

1 Motivation and Scope

After the Basel II accord in 2004, ratings [1] became an increasingly important instrument in Credit Risk [2], as they allow banks to base their capital requirements on internal as well as external rating systems. These ratings became instrumental in evaluating the risk of a bond or loan and in the calculation of the Value at Risk. As such, it is often desirable to quantify the uncertainty in these ratings, and predict the likelihood that an institution will be up- or down-graded in the near future. A common technique is to aggregate credit rating transition data over yearly or quarterly periods, and to model future transitions using these data. However, to be reliably used in this way the ratings' evolution must obey particular features that can be evaluated through analysis of the data published by rating agencies. Two important properties that can be taken as sufficient for accepting the empirical data as a reliable indicator of future rating evolution are time homogeneity and Markovianity.

The representation of the evolution of a time continuous process by an aggregated transition matrix will not be adequate if the underlying process is not Markov, or not time homogeneous [3]. Furthermore, different techniques to estimate a transition matrix from a finite sample of data should be employed depending on whether the process is time homogeneous or not [4]. Theoretically, both the Markov and the time homogeneous assumptions simplify considerably the models in question [5], but typically only the latter is at times dropped in order to build a more general theoretical framework.

In this paper we test how good both assumptions are in different periods of time for a homogeneous rating class in Moody's database. We compare transition matrices calculated under different assumptions and show that the quality of the time homogeneous and Markov assumptions change considerably in time. Moreover, we argue that the wide fluctuations of the assumptions' quality may on the one hand provide evidence for detecting discontinuities in the rating process, e.g. when establishing new evaluation criteria for a bank rating, and, on the other hand, can be taken as a tool for ascertaining how complete and trustable such rating criteria are.

We start in Sec. 2 describing the empirical data collected from Moody's and in Sec. 3 we describe how to test the validity of both the homogeneity and Markovianity assumptions, Secs. 3.1 and 3.2 respectively. Section 4 concludes the paper and presents some discussion of our results in the light of finance rating procedures.

2 Data: Ten Years of Rating Transitions in Europe

The data analyzed in this paper was collected directly from Moody's website [6]. We used a partially automated procedure to extract the data and build the rating time series of each bank.

The rating time series of each bank has a sample frequency of one day, starts in January 1st 2003 and ends in January 1st of 2013. The data sample is the set of rating histories from the banks, in European countries, that had a rating at the final date. Each value indicates the rating class, according to the so-called *Banking Financial Strength* [7], at which the bank is evaluated at that particular day.

One first important feature of this rating database is its non-stationary character, as can be seen in Fig. 1. The number of banks N_R included in the data set increased almost monotonically during the total time-span analyzed by us (see Fig. 1a). On January 1st 2003 there are $N_R = 446$ rated companies in the data set, and this number increases until 2013 when one registers $N_R = 924$ rated banks. Therefore, we will consider our measures normalized to the number of banks in the database. Further, we count in Moody's database a total of $N_T = 1432$ rating transitions, that distribute heterogeneously in time. Indeed, the number of transitions N_T per bank also changes significantly, with three events of peaked activity, namely during

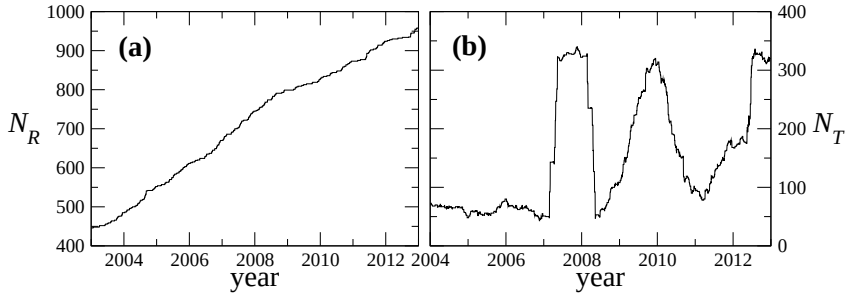

Fig. 1. (a) Number of bank entities in Moody's data sample as a function of time and **(b)** the number of transitions per bank, computed as moving averages during one-year periods.

the year of 2007, at the beginning of 2010 and in the last half year of 2012 (see Fig. 1b). This will be of importance when analyzing the evolution of the generator homogeneity and Markovianity of the corresponding transition matrices.

The rating category is a measure of the capacity of the institution to meet its financial obligations and avoid default or government bailout. We have $n_s = 15$ rating states, denoted by characters A to E in alphabetic order and with the two possible extra suffixes, namely $+$ and $-$. State $A+$ represents the state corresponding to the best financial health and less credit risk, followed by A, $A-$, $B+$ and so on, until the bottom of the scale, $E-$, the state that represents the highest risk level. Figure 2 shows three plots (left) illustrating the histogram of rating states at three different time, namely the first day of 2007, 2009 and 2010.

Henceforth, we define $\tilde{R}_i(t)$ as the rating of the bank number i at the moment t, and we map the rating states to an increasing ordered number series: state $\tilde{R} = E+$ corresponding to label $R = 0$, and state $\tilde{R} = A+$ to label $R = 14$. With such a labelling it is possible to compute rating increments as

$$T_i(t, \tau) = R_i(t) - R_i(t - \tau). \tag{1}$$

When $T_i(t) > 0$ (resp. < 0) it means that bank i saw its rating increased (resp. decreased) at time t. Unless stated otherwise we will use always $\tau = 365$ days. The plots in the right column of Fig. 2 show the histograms of the corresponding rating increments at the same three days.

We call henceforth $R(t)$ and $T(t)$ the aggregated processes of the ratings and rating increments respectively, over all N_R companies observed at time t. Figure 3 shows the evolution of the first four moments for both rating distributions (left) and transition distributions (right), with $\tau = 365$ days.

The average rating $\langle R \rangle$ (Fig. 3a) has decreased during most of the ten year period records. We should note however that this is due to the new entries in

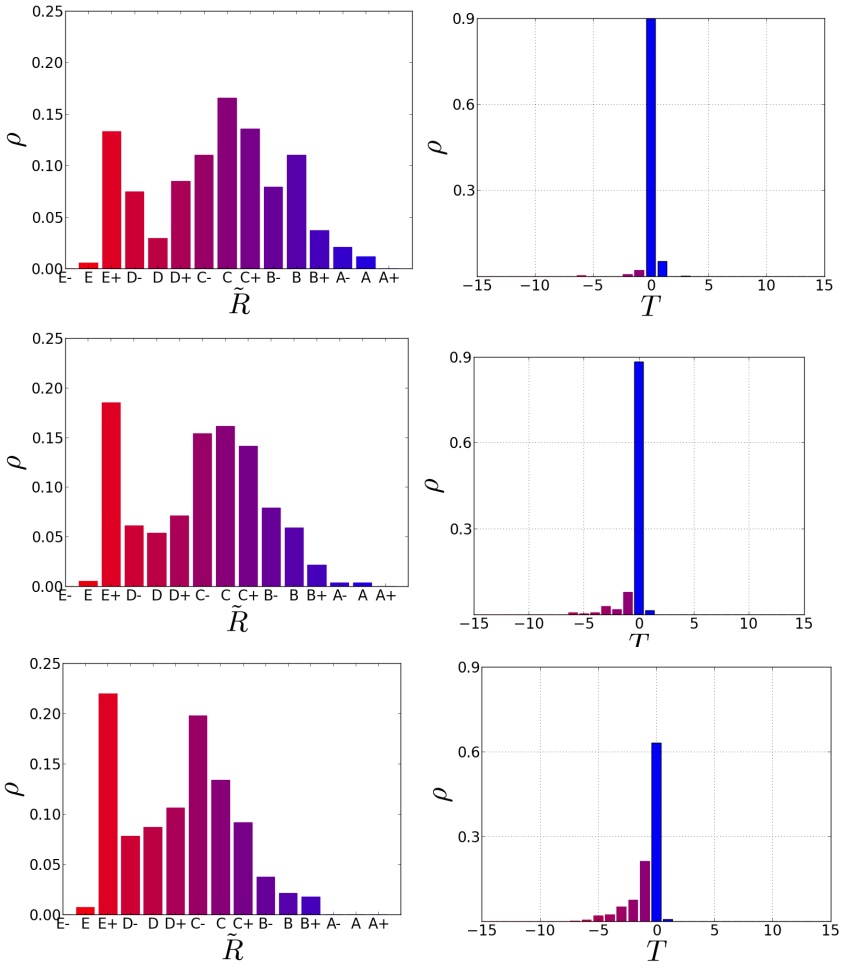

Fig. 2. Illustration of rating histograms for the rating state $\tilde{R}$ (left) and the corresponding rating variations $T = \Delta R$ (right), where R is an integer enconding the rating state, ranging from 0 ($E-$) to 14 ($A+$). Three different days are selected: first day of 2007 (first row), 2009 (second row) and 2010 (third row); cf. Fig. 3.

the database with an initial low rating, since $\langle T \rangle$ has positive periods during the first five years of the recorded set.

As for the rating variance σ_R (Fig. 3e), after a slight increase, it also decreased since 2007, due to the concentration of rating states to the lower rating classes ($\langle T \rangle < 0$). The transitions however exhibit two periods of increased variance σ_T (Fig. 3f), which reflect probably the respective increase in the number of transitions (compare with Fig. 1b).

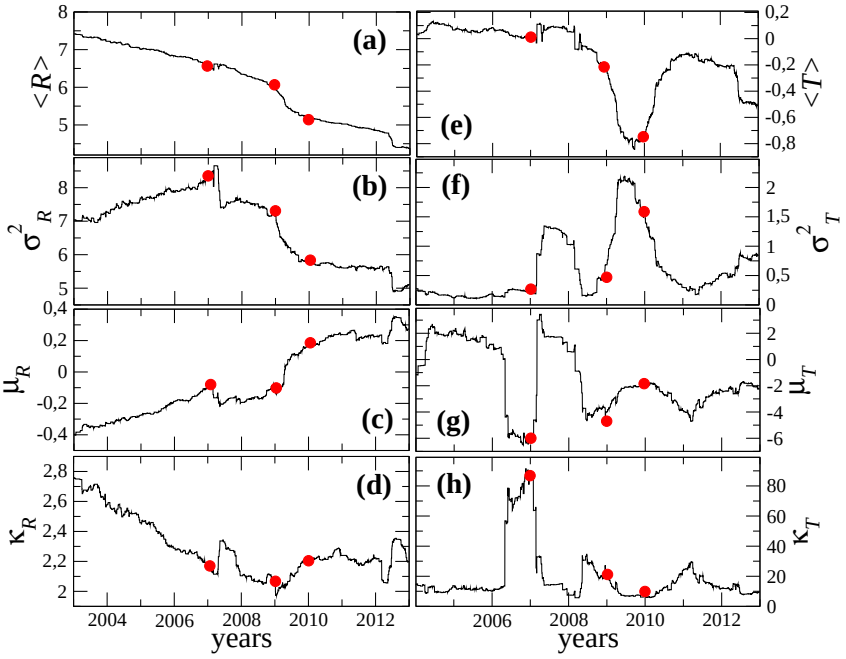

Fig. 3. Evolution of the first statistical moments of **(a-d)** the rating state R distribution and **(e-h)** its 1year-increment T distribution. From top to bottom: averages $\langle R\rangle$ and $\langle T\rangle$, variances σ_R^2 and σ_T^2, skewnesses μ_R and μ_T, kurtosis κ_R and κ_T. Red bullets indicate the days when the histograms in Fig. 2 were taken.

As the lowest states get more and more dominant, the rating skewness μ_R (Fig. 3c) increases steadily, until it changes sign around 2008, when transitions become negative on average. These two observations are consistent with each other: the negative skewness indicates the large majority of banks being below the average rating which corresponds to an average decrease of the rating $\langle T\rangle < 0$. It also indicates that there are a few banks highly rated. This observation together with the observations regarding temporal homogeneity in the next section will justify some comments about the objectiveness of rating criteria.

The rating distribution is also typically platykurtic (see Fig. 3d), as its kurtosis is always below three (Gaussian kurtosis), indicating a more pronounced flatness around the average of rating distributions. Concerning the third and fourth moments of transition distributions, Figs. 3g and 3h respectively, we see large fluctuations during the periods with fewer transitions. One can clearly see a very high kurtosis, and changes in the sign of the mean and skewness.

3 What is the Underlying Continuous Process?

It will be assumed in the following that the process governing rating transitions has a continuous processes underlying it, an assumption which has been the subject of previous investigation without a clear result, see e.g. [8]. Even in case that there is a continuous process, the corresponding generator may be constant (homogeneous generator) or vary in time (non-homogeneous).

The non-homogeneity is important in the finance context since it limits the range of models that can be used. In particular, it has been argued [3] that a better method of estimating a transition matrix than the usual one exists, if we consider time-homogeneity. The main advantages of this method are capturing very small transition probabilities between two states, even when no transitions occurred between those two states, and the distinction on whether a transition occurred earlier or later within the studied time-frame. The time-homogeneity condition is also important to check if the rating philosophies [9,10] allegedly used are being correctly followed or not. This assumption is also not expected to hold if criteria by which ratings are ascribed to banks are not constant in time, but vary according to artificial or externally imposed factors [11].

Furthermore, another important feature of continuous transition processes is their Markovianity. The Markov property is important if the current rating of a bank is to be considered a complete indicator of its future risk. In this section we will address both these conditions separately.

3.1 Testing Time-Homogeneity

Mathematically, if a time continuous Markov process is time homogeneous then there is a constant matrix $\mathbf{Q}$, called a generator, solution of

$$\frac{d\mathbf{M}(t)}{dt} = \mathbf{Q}\mathbf{M}(t), \tag{2}$$

where $\mathbf{M}$ is the transition matrix, with entries M_{ij} given the probability for observing a transition from state i to state j $(i, j = 1, \ldots, n_s)$. In other words, a time continuous process is time homogeneous if, being Markov, its transition matrix can be expressed as $\mathbf{M}(t) = e^{\mathbf{Q}t}$, and therefore it has a well-defined logarithm. We take the analogue from ordinary differential equations and loosely call $\mathbf{Q}$ the logarithm of $\mathbf{M}$.

The mathematical conditions for the existence of a homogeneous generator give a bivalent result [8,12] that does not take into consideration noise generated from finite samples, nor how distant an empirical process is from being time continuous. Therefore, we neglect several mathematical results that determine if a generator exists or not, and assume that the process is Markov and continuous. Being Markov and continuous means that there is a generator satisfying Eq. (2) and that is either constant or varies in time.

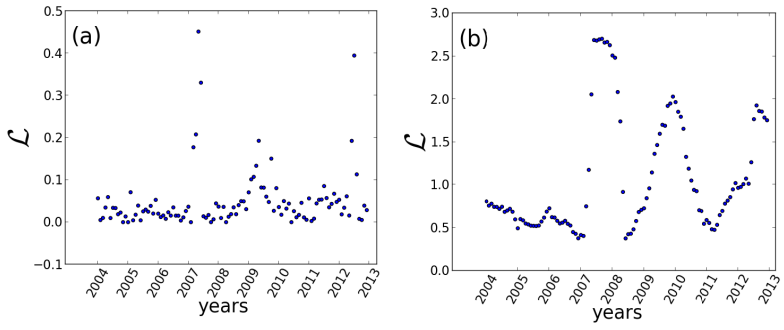

Fig. 4. Testing for temporal homogeneity: difference between the log-likelihood $\mathcal{L}$ of the transition matrix $\mathbf{M}^{(e)}$ and the transition matrix $\mathbf{M}$ calculated assuming time homogeneity. Both matrices are calculated over a time interval **(a)** $\tau = 1$ month and **(b)** $\tau = 1$ year. The log-likelihood Eq. (4) difference was calculated at the first day of each month from January 2004 to December 2012.

Next, we estimate the closest constant generator $\mathbf{Q}$ directly from the empirical data, compute the associated matrix $\mathbf{M} = e^{\mathbf{Q}t}$, and compare it with the empirical transition matrix $\mathbf{M}^{(e)}$. For estimating the generator matrix $\mathbf{Q}$ we follow the approach described in Ref. [4], calculating its off-diagonal elements as

$$Q_{ij} = \frac{N_T^{(ij)}}{\int_{t_0}^{t_f} N_R^{(i)}(t)dt}, \tag{3}$$

where $N_T^{(ij)}$ represents the number of transitions from i to j between the times t_0 and t_f. The number $N_R^{(i)}(t)$ stands for the number of banks in state i at time t. The diagonal elements Q_{ii} follow from the condition $\sum_j Q_{ij} = 0$.

To compute the distance between a time homogeneous process and the empirical process we compare $\mathbf{M}$ with $\mathbf{M}^{(e)}$, and plot the statistic:

$$\mathcal{L} = \frac{\sum_{i,j} N_T^{(ij)} \left(\log M_{ij} - \log M_{ij}^{(e)} \right)}{\sum_{i,j} N_T^{(ij)}} \tag{4}$$

This is a log-likelihood ratio; loosely speaking it quantifies the error introduced by making the more restrictive assumption of time homogeneity. The results are shown in Fig. 4: in panel (a) we aggregate the data in periods of one month while in panel (b) the aggregation period is one year.

It can be seen that there are three periods when the time homogeneity condition becomes an insufficient approximation to the dynamics of the process. The first period starts in the early 2007, the second period around the middle of 2009, and the third period in the last half of 2012. The profile of the time inhomogeneity is different for each time period. It shows a sharp peak in 2007, concentrated in just a few months, and wider in the other periods.

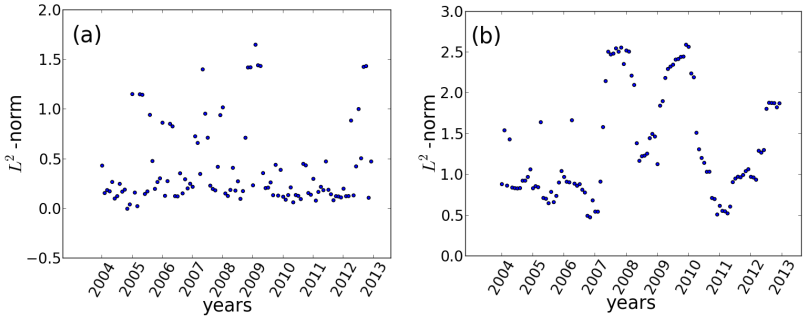

Fig. 5. Testing Markovianity: difference between the empirical transition matrix $\mathbf{M}_{0\tau}^{(e)}$ calculated over a time-interval $[0, \tau]$ and the product of the half-period matrices, $\mathbf{M}_{0\frac{\tau}{2}}^{(e)}$ and $\mathbf{M}_{\frac{\tau}{2}\tau}^{(e)}$, using the L^2-norm defined in Eqs. (7,8). **(a)** $\tau = 1$ month and **(b)** $\tau = 1$ year. The difference was calculated at the first day of each month between January 2004 and December 2012.

It is interesting to analyse these periods in the previous Figs. 1-3: early 2007, 2009 and 2010. In 2007 there was an unusually high number of rating transitions, even considering that only about 700 companies were rated at the time. In Fig. 3 it can be seen that in this period the variance σ_R of the ratings decreased, the skewness μ_R had slight negative burst, and there was an increase in the kurtosis κ_R. As for the statistics of transitions, one can see in this period that the average $\langle T \rangle$ becoming positive, the skewness (μ_T) changing signal and becomes positive and the kurtosis (κ_T) becomes much smaller. The variance of T increases, but again that can be explained by the high number of rating transitions in that period.

In late 2009 and early 2010 we have a very different profile. In this period the downgrades are the rule, as one can see by the negative values of $\langle T \rangle$. The relatively low values of κ_T and the absolute value of μ_T tells us that this was a general trend, and not a very drastic movement by just a few banks.

In 2012 the scenario is similar. Again there are more downgrades, and this a general trend. The companies are now much more condensed, as one can see by the low values in σ_R in the bad quality part of the rating scale around $\langle R \rangle$.

3.2 Testing the Markov Hypothesis

Mathematically, a Markov process x_t obeys the following condition:

$$\Pr(x_{t_1}|x_{t_2}, x_{t_3}, \dots) = \Pr(x_{t_1}|x_{t_2}) \tag{5}$$

with $t_1 > t_2 > t_3 > \dots$. The conditional probability in the right hand-side of Eq. (5), $\Pr(x_{t_1}|x_{t_2})$, is exactly specified by the transition matrix $\mathbf{M}$.

The rating process must be assumed to be Markov, otherwise a rating would not represent a uniform risk class, as its elements could be distinguished according to their previous series of rating states.

From the definition of a Markov process in Eq. (5) it is straightforward to show that a Markov process also obeys

$$\mathbf{M}_{t_0 t_f} = \prod_{n=1}^{N} \mathbf{M}_{t_{n-1} t_n}\,, \tag{6}$$

where N is the number of subintervals in $[t_0, t_f]$ and labels $t_i t_j$ denote the time interval $[t_i, t_j]$ considered when determining $\mathbf{M}_{t_i t_j}$. Here we fix $N = 2$ and consider two equally spaced intervals with $\tau \equiv t_f - t_0 = 1$ month and $\tau = 1$ year. Equation (6) is known as the Chapman-Kolmogorov equation [13] and it does not hold in general either when the process is non-Markov or when we have an insufficiently short sample of data.

We will use the Chapman-Kolmogorov equation as a test indicating whether the rating database of Moody's is Markov. To that end, we consider empirical matrices $\mathbf{M}^{(e)}_{0\tau}$ computed for full 1 month or 1 year intervals, and compare it with the associated product of the two corresponding half-periods, $\overline{\mathbf{M}^{(e)}}_{0\tau} = \mathbf{M}^{(e)}_{0\frac{\tau}{2}} \mathbf{M}^{(e)}_{\frac{\tau}{2}\tau}$. For the comparison we now use the L_2-norm instead of the the $\mathcal{L}$ log-likelihood, since the latter creates singularities when dealing with zero entries that are now more frequent. The L2-norm of a square matrix $\mathbf{A}$ is defined through the maximum singular value of $\mathbf{A}$,

$$\|\mathbf{A}\| = \sigma_{\max}(\mathbf{A})\,, \tag{7}$$

and we compute it for the difference

$$\mathbf{A} = \mathbf{M}^{(e)}_{0\tau} - \overline{\mathbf{M}^{(e)}}_{0\tau}\,, \tag{8}$$

where $\|\cdot\|$ represents the usual Euclidian norm.

Results are shown in Fig. 5. Clearly, there are two periods when the Markov assumption seems less valid. The first period is in early 2007, and the second in the middle of 2009, followed by another, less significant increase at the end of 2012. As said before, this coincides with an abrupt change in the statistics of T and R.

4 Discussion and conclusions

We have addressed time series of credit ratings provided by Moody's and studied simple ways to compute the validity of the time homogeneous and Markov assumptions. We have shown how the accuracies of the time homogeneous assumption and Markov assumption vary with time. Naturally, when the Markov assumption fails, so does the time homogeneous one, in

particular during 2007 and in the latest half of 2009 and beginning of 2010. In these periods the statistics of the process changed considerably. In the end of the year of 2012 the accuracy of the time homogeneous assumption is low but the Markov approximation is within the usual fluctuation range. In this period there is a less abrupt change in the statistics of the process.

One must stress that when the Markov assumption does not hold, the ratings are not a complete measure of the risk of a given entity, since further information besides the actual rating needs to be specified. Moreover, our results raise the evidence that perhaps in 2007 new rating criteria were introduced, imposing a discontinuity in the series of ratings, or that new rating transition were correlated with previous ones, which could support the claim that rating agencies were an active part in the crisis that followed.

Our approach can be improved by introducing for instance a more sophisticated procedure for extracting the histograms for the ratings and their increments, namely using the kernel based density, which is known to converge faster to the real distribution than the usual binning procedure. From this first approach to investigate Moody's rating database one can now attack the embedding problem for the series of transition matrices, where different generators estimates can be compared. These and other issues will be addressed elsewhere.

Acknowledgments

The authors thank Fundação para a Ciência e a Tecnologia for financial support under PEst-OE/FIS/UI0618/2011, PEst-OE/MAT/UI0152/2011, FCOMP-01-0124-FEDER-016080, SFRH/BPD/65427/2009 (FR). This work is part of a bilateral cooperation DRI/DAAD/1208/2013 supported by FCT and Deutscher Akademischer Auslandsdienst (DAAD).

References

1. Robert A Jarrow, David Lando, and Stuart M Turnbull. A markov model for the term structure of credit risk spreads. *Review of Financial studies*, 10(2):481–523, 1997.
2. Basel Committee on Banking Supervision. Basel II: International convergence of capital measurement and capital standards: a revised framework. *(available at http://bis.org/publ/bcbs107.htm)*, 2004.
3. T. Charitos, P.R. de Waal, and L. C. van der Gaag. Computing short-interval transition matrices of a discrete-time markov chain from partially observed data. *Statistics in medicine*, 27(6):905–921, 2008.
4. D. Lando and T. M. Skødeberg. Analyzing rating transitions and rating drift with continuous observations. *Journal of Banking & Finance*, 26(2):423–444, 2002.
5. R. Weißbach, P. Tschiersch, and C. Lawrenz. Testing time-homogeneity of rating transitions after origination of debt. *Empirical Economics*, 36(3):575–596, 2009.

6. https://www.moodys.com/researchandratings.
7. Moody's Investors Service. Moody's rating symbols & definitions. *Report*, 79004(08):1–52, 2004.
8. R.B. Israel, J.S. Rosenthal, and J.Z. Wei. Finding generators for markov chains via empirical transition matrices, with applications to credit ratings. *Mathematical Finance*, 11(2):245–265, 2001.
9. J. Kiff, M. Kisser, and L. Schumacher. An inspection of the through-the-cycle rating methodology. *IMF Working Paper*, 2013.
10. Z. Varsanyi. Rating philosophies: some clarifications. *Report*, 2007.
11. J. Mathis, J. McAndrews, and J.-C. Rochet. Rating the raters: are reputation concerns powerful enough to discipline rating agencies? *Journal of Monetary Economics*, 56(5):657–674, 2009.
12. EB Davies. Embeddable markov matrices. *Electronic Journal of Probability*, 15:1474–1486, 2010.
13. H. Risken. *The Fokker-Planck Equation*. Springer, Berlin, 2nd edition, 1989.